

# **PERIODIC CALIBRATION FOR COMMUNICATION CHANNELS BY DRIFT TRACKING**

Inventors: Craig E. Hampel  
Frederick A. Ware  
Richard E. Perego

## **BACKGROUND OF THE INVENTION**

### Field of the Invention

[0001] The present invention relates to the calibration of communication channel parameters in systems, including mesochronous systems, in which two (or more) components communicate via an interconnection link; and to the calibration needed to account for drift of conditions related to such parameters during operation of the communication channels.

### Description of Related Art

[0002] In mesochronous communication channels, typically a reference clock provides frequency and phase information to the two components at either end of the link. A transmitter on one component and a receiver on another component each connect to the link. The transmitter and receiver operate in different clock domains, which have an arbitrary (but fixed) phase relationship to the reference clock. The phase relationship between transmitter and receiver is chosen so that the propagation delay seen by a signal wavefront passing from the transmitter to the receiver will not contribute to the timing budget for signaling rate. Instead, the signaling rate is determined primarily by the drive window of the transmitter and the sample window of the receiver. The signaling rate will also be affected by a variety of second order effects. This system is clocked in a mesochronous fashion, with the components locked to specific phases relative to the reference clock, and with the drive timing point and sample timing point of each link fixed to the phase values that maximize the signaling rate.

[0003] These fixed phase values may be determined in a number of ways. A sideband link may accompany a data link (or links), permitting phase information to be passed between transmitter and receiver. Alternatively, an initialization process may be invoked when the system is first given power, and the proper phase values determined by setting a transmitter drive timing point, and executing a calibration sequence that includes

passing a calibration pattern or calibration patterns across the actual link to the receiver, and adjusting the receiver sample timing point until the data is successfully sampled at the receiver. Once the drive timing point and sample timing point of each link have been fixed, the system is permitted to start normal operations. The calibration sequences used to establish proper phase values on initialization are designed to provide reliable results over a wide variety of environmental conditions, and a wide variety of operating conditions on the communication link. To provide reliable results, the calibration sequences pass a calibration pattern or calibration patterns that include lots of data and take quite a long time.

**[0004]** The calibration sequences use patterns with long or numerous codes designed to find the worst-case leading edge and the worst-case trailing edge of a passing region for a parameter subject of the calibration, such as drive timing points and sample timing points discussed above. The edge values are a function of many system parameters including for example silicon processing variations, and packaging parameters such as crosstalk, terminal resistance accuracy, system board impedance, module impedance, trace lengths, connector locations, and so on. In addition to the uncertainty of these variables, the patterns necessary to create worst-case inter-symbol interference or resonance can be very long, and are difficult to predict. In order to deal with the uncertainty of which patterns will generate the leading or trailing edge of the worst-case passing region, many systems use a brute force approach to calibration sequences, using sequences with very long calibration patterns. For example, one brute force approach is based on the use of a pseudorandom bit sequence PRBS, which consists of a long, fairly random pattern, to attempt to determine the passing region. Other systems utilize many initialization patterns that are hundreds of bits long to present the worst-case pattern for a given configuration, and attempt to cover all possible conditions.

**[0005]** The use of long, complex patterns is generally adequate if the calibration sequence is run infrequently. For example, if the algorithm is run only during an initial system bring up, the length of the calibration pattern is not generally critical. However, during normal operation, system conditions will change. Ambient temperature, component temperature, supply voltages, and reference voltages will drift from their initial values. Also, spread spectrum clock systems intentionally shift the clock

frequency to meet emission standards. As the conditions drift, the optimal timing points of the transmitter and receiver and other parameters will change.

[0006] Although the calibration sequences can be run periodically to adjust for drift, the length and complexity of the patterns and the algorithms used in the sequences interfere with mission-critical operations of the system. Primarily, the long calibration sequences with patterns having long or numerous codes require storage, access to the input/output circuits and interconnect, and processing resources in order to complete a calibration sequence. During this recalibration time, the system is unavailable to the application. This creates at least two significant problems. First, performance is reduced, in general. Second, many applications can tolerate only minimum latency addition without under or over running their streaming data.

[0007] It is desirable to provide techniques to compensate for the condition drift, and provide improvements in system and component design to permit these techniques to be utilized.

## SUMMARY OF THE INVENTION

[0008] The present invention takes advantage of a discovery that long-term drift in timing, voltage or other parameters of a communication channel caused for example by temperature change, can be largely independent of the calibration pattern utilized to compute the operation value of the parameter. In other words, drift in a particular parameter can be tracked without using patterns based on long and/or numerous codes required for determining an optimal operation value for the parameter. Rather, a simpler calibration sequence can be applied to track drift of the particular parameter during operation of the device. The simpler calibration sequence used to track drift of the parameter utilizes less of the resources of the communication channel (that is, it occupies the link, transmitter and receiver for less time) than the more exhaustive initial calibration sequence. Changes in the parameter based on computations that arise from the use of a simpler calibration sequence are used to adjust the operation value that was created using the more exhaustive initial sequence. Herein, the operation value is the value, or a function of the value, of the parameter usually used during normal operation of the communication channel.

[0009] Accordingly, the present invention provides a system and method for calibrating an operation value for a parameter of a communication channel, which allows for optimizing the operation value of the parameter, while efficiently accounting for drift of properties of the channel. A method for calibrating a communication channel,

5 including a first component having a transmitter, a second component having a receiver, and a communication link coupling the first and second components, the communication channel having a parameter with an operation value determined by calibration, comprises establishing an operation value for a parameter of the communication channel; executing a drift calibration sequence, from time to time, to determine a drift value for the

10 parameter of the communication channel, wherein drift calibration sequence comprises an algorithm different than used to establish the operation value; and updating the operation value in response to the drift value. In embodiments of the invention, the operation value is established by executing a first calibration sequence to set the operation value of the parameter of the communication channel, and wherein the drift

15 calibration sequence utilizes less resources of the communication channel than the first calibration sequence. The first calibration sequence is executed, such as upon initialization of a system, to establish the operation value, which utilizes an algorithm that the designer intends to be exhaustive so that a suitable operation value can be determined for most conditions in which the system is designed to operate. A second calibration

20 sequence is executed to determine a drift value in the parameter, and to update the operation value in response to the drift value. The second calibration sequence is executed from time to time during normal operation of the communication channel, and utilizes less resources of the communication channel than does the first calibration sequence. In one embodiment, the first calibration sequence for measurement of the

25 operation value utilizes a long calibration pattern, such a data set including a code or codes which includes more than 200 bits, a code greater than 30 bytes, and a code based on a pseudorandom bit sequence having a length of  $2^N - 1$  bits, where N is equal to or greater than 7. The second calibration sequence utilizes a short calibration pattern, such as a fixed code less than 130 bits, such as less than or equal to 16 bytes, and for example

30 as short as 2 bytes long. Thus in embodiments of the invention, the first calibration sequence utilizes many relatively short codes, or combinations of relatively short and

longer codes, while the second calibration sequence uses one short code, or a few short codes where the total number of bits used in the second sequence is substantially less than the total number of bits used in the first sequence.

[0010] A calibration sequence, in some embodiments, includes a plurality of

5 calibration cycles. The calibration cycles include de-coupling, logically using software or physically using switches, the normal signal source from the transmitter and supplying a calibration pattern in its place. The calibration pattern is transmitted on the link using the transmitter on the first component. After transmitting the calibration pattern, the normal signal source is re-coupled to the transmitter. The calibration pattern is received  
10 from the communication link using the receiver on the second component. The received calibration pattern is analyzed, by comparing it to an expected calibration pattern. The comparison indicates a number of errors in the received calibration pattern, which are used to indicate whether the value of the parameter used during the calibration sequence allows successful operation of the communication channel. A calibrated value of a  
15 parameter, or drift of a previously established parameter, of the communication channel is determined in response to the analysis of the received calibration pattern or patterns.

[0011] Some embodiments of the invention comprise a calibration method, based on the first and second calibration sequences as discussed above, wherein said second calibration sequence includes

- 20 adjusting the parameter to a first edge value;
- transmitting calibration patterns on the communication link
- using the transmitter on the first component;
- receiving the calibration patterns on the communication link
- using the receiver on the second component;
- 25 determining whether the received calibration patterns indicate
- that a new edge value is met, and if not, returning to the adjusting to
- iteratively execute a next calibration cycle, and if a new edge value
- is met, then saving data indicating drift of the first edge value; and
- then
- 30 adjusting the parameter to a second edge value;

transmitting calibration patterns on the communication link  
using the transmitter on the first component;

receiving the calibration patterns on the communication link  
using the receiver on the second component;

5 determining whether the received calibration patterns indicate  
that a new edge value is met, and if not, returning to the adjusting to  
iteratively execute a next calibration cycle, and if a new edge value  
is met, then saving data indicating drift of the second edge value;  
and

10 determining the drift for the parameter in response to a function  
of the drift of the first and second edge values.

**[0012]** In some embodiments of the invention, the second calibration sequence  
measures only one of the first and second edge values, or other value related to the  
operation value, and determines the drift for the parameter in response to a function of the  
15 one measured value.

**[0013]** Methods according to some embodiments of the invention comprise  
executing a first calibration sequence;  
the first calibration sequence including iteratively adjusting a  
value of a first edge parameter, transmitting a long calibration  
20 pattern using the transmitter on the first component, receiving the  
long calibration pattern using the receiver on the second component,  
and analyzing the received long calibration pattern to determine a  
value for the first edge parameter;

the first calibration sequence also including iteratively adjusting  
25 a value of a second edge parameter, transmitting a long calibration  
pattern using the transmitter on the first component, receiving the  
long calibration pattern using the receiver on the second component,  
and analyzing the received long calibration pattern to determine a  
value for the second edge parameter; and

30 as a result of the first calibration sequence, determining said  
operation value for the parameter based on the function of the first

and second edge parameters;

executing a second calibration sequence;

the second calibration sequence including iteratively adjusting a value of a first edge parameter, transmitting a short calibration pattern using the transmitter on the first component, receiving the short calibration pattern using the receiver on the second component, and analyzing the received short calibration pattern to determine a value for the first edge value;

the second calibration sequence also including iteratively adjusting a value of a second edge parameter, transmitting a short calibration pattern using the transmitter on the first component, receiving the short calibration pattern using the receiver on the second component, and analyzing the received short calibration pattern to determine a value for the second edge value; and

as a result of the second calibration sequence, determining said drift for the parameter based on the function of the first and second edge parameters; and

updating said operation value based on said drift, wherein the long calibration patterns have lengths of at least 30 bytes, and the short calibration patterns have lengths of 16 bytes, or less.

**[0014]** Other aspects and advantages of the present invention can be seen on review of the drawings, the detailed description and the claims, which follow.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0015]** Fig. 1 is a simplified diagram of two components interconnected by a communication channel.

**[0016]** Fig. 2 is a timing diagram illustrating timing parameters for a communication channel like that shown in Fig. 1.

**[0017]** Fig. 3 illustrates an embodiment of the present invention where both a transmitter drive point and a receiver sample point are adjustable.

[0018] Fig. 4 illustrates an embodiment of the present invention where only a receiver sample point is adjustable.

[0019] Fig. 5 illustrates an embodiment of the present invention where only a transmitter drive point is adjustable.

5 [0020] Fig. 6 is a chart illustrating measurement of an operation value and measurement of drift for a transmit clock parameter TX of a communication channel using long and short calibration patterns.

[0021] Fig. 7 is a chart illustrating measurement of an operation value and measurement of drift for a receive clock parameter RX of a communication channel using  
10 long and short calibration patterns.

[0022] Fig. 8 is a flow chart illustrating calibration steps for a transmitter on a unidirectional link for a transmitter drive point.

[0023] Fig. 9 illustrates timing for iteration steps for calibrating a transmitter drive point.

15 [0024] Fig. 10 is a flow chart illustrating calibration steps for a receiver on a unidirectional link for a sample point.

[0025] Fig. 11 illustrates timing for iteration steps for calibrating a receiver sample point.

[0026] Fig. 12 illustrates an embodiment of the present invention where transmitter drive points and receiver sample points on components of a bidirectional link are  
20 adjustable.

[0027] Fig. 13 illustrates an embodiment of the present invention where receiver sample points on components of a bidirectional link are adjustable.

[0028] Fig. 14 illustrates an embodiment of the present invention where both  
25 components have adjustable transmitter drive points.

[0029] Fig. 15 illustrates an embodiment of the present invention where a transmitter drive point and a receiver sample point of only one component on a bidirectional link are adjustable.

[0030] Fig. 16 is a flow chart illustrating calibration steps for a transmitter drive point  
30 for a bidirectional link.



[0031] Fig. 17 is a flow chart illustrating calibration steps for a receiver sample point for a bidirectional link.

[0032] Fig. 18 illustrates an embodiment of the present invention where a transmitter drive point and a receiver sample point of one component on a bidirectional link are adjustable with a plurality of parameter sets, and wherein the bidirectional link is coupled to a plurality of other components corresponding to the plurality of parameter sets.

## DETAILED DESCRIPTION

[0033] A detailed description of embodiments of the present invention is provided with reference to the Figures.

### Transmitter and Receiver Timing Parameters

[0034] Fig. 1 shows two components 10, 11 connected with an interconnection medium, referred to as Link 12. One has a transmitter circuit 13 which drives symbols (bits) on Link 12 typically in response to rising-edge timing events on the internal CLK<sub>T</sub> signal 14. This series of bits forms signal DATAT. The other component 11 has a receiver circuit 15 which samples symbols (bits) on Link 12 typically in response to rising-edge timing events on the internal CLK<sub>R</sub> signal 16. This series of bits forms signal DATAR. The DATAT and DATAR signals are related; DATAR is an attenuated, time-delayed copy of DATAT. The attenuation and time-delay occur as the signal wavefronts propagate along the interconnection medium of Link 12. The communication channel comprising transmitter 13, link 12 and receiver 15 can be characterized by a plurality of parameters that have an operation value during normal operation of the channel. Representative parameters of the channel include timing parameters specifying the drive times and sample times, voltage levels for drivers and comparators, resistance values such as link termination resistances, driver strength, adaptive equalization coefficients, noise cancellation coefficients, parameters that cause overshoot and undershoot of signals such as driver switching power or speed, and so on. These parameters are affected by or contribute to, drifting flight times, output delays for transmitters, receiver mismatches, input delays, and so on.

[0035] Fig. 2 illustrates timing parameters for a mesosynchronous channel, including the transmit clock CLK<sub>T</sub> signal 14 on trace 20, the transmitter signal DATAT on trace

21, the receive clock CLKR signal 16 on trace 22, and the receiver signal DATAR on trace 23. The transmitter eye 24 and the receiver eye 25 are also illustrated.

[0036] The transmitter circuit 13 will begin driving a bit (labeled “a”) no later than a

time  $t_{Q,MAX}$  after a rising edge 30 of CLKT, and will continue to drive it until at least a

5 time  $t_{V,MIN}$  after the next rising edge 31.  $t_{Q,MAX}$  and  $t_{V,MIN}$  are the primary timing parameters of the transmitter circuit 13. These two values are specified across the full range of operating conditions and processing conditions of the communication channel.

As a result,  $t_{Q,MAX}$  will be larger than  $t_{V,MIN}$ , and the difference will represent the dead time or dead band 32 of the transmitter circuit 13. The transmitter dead band 32

10 ( $t_{DEAD,T}$ ) is the portion of the bit timing window (also called bit time or bit window) that is consumed by the transmitter circuit 13:

$$t_{DEAD,T} = t_{Q,MAX} - t_{V,MIN}$$

[0037] The edges of the timing window around transition 31 can be defined by:

$$t^{-1} + t_{Q,MAX}, \text{ and}$$

15  $t + t_{V,MIN},$

where  $t^{-1}$  is transition 30 one clock cycle earlier than the transition 31, and  $t$  is the transition 31.

[0038] The receiver circuit 15 will begin sampling a bit (labeled “a”) no earlier than a

20 time  $t_{S,MIN}$  before a rising edge 35 (or 36) of CLKR, and will continue to sample until no later than a time  $t_{H,MIN}$  after the rising edge 35.  $t_{S,MIN}$  and  $t_{H,MIN}$  are the primary timing parameters of the receiver circuit. These two values are specified across the full range of operating conditions and processing conditions of the circuit. The sum of  $t_{S,MIN}$  and

$t_{H,MIN}$  will represent the dead time or dead band 37, 38 of the receiver. The receiver dead band 37, 38 ( $t_{DEAD,R}$ ) is the portion of the bit timing window (also called bit time or bit window) that is consumed by the receiver circuit:

$$t_{DEAD,R} = t_{S,MIN} + t_{H,MIN}$$

[0039] The edges of the timing window around transition 35 can be defined by:

$$t - t_{S,MIN}, \text{ and}$$

30  $t + t_{H,MIN},$

where  $t$  is transition 35.

[0040] In this example, the bit timing window is one  $t_{\text{CYCLE}}$  minus the  $t_{\text{DEAD,T}}$  and  $t_{\text{DEAD,R}}$  values, each of which is about 1/3 of one  $t_{\text{CYCLE}}$  in this example. The remaining 1/3  $t_{\text{CYCLE}}$  would account for other uncertainty. Such uncertainty could include, for example, variation in the placement of the rising edges of CLKT and CLKR. In some systems, this variation might be specified as part of the  $t_{\text{DEAD,T}}$  and  $t_{\text{DEAD,R}}$  definition. Other uncertainty could include variation in the propagation delay across the interconnection medium.

### Unidirectional Link Alternatives

[0041] Fig. 3 shows two components 100 (transmit component) and 101 (receive component) connected with an interconnection medium referred to as Link 102. The link is assumed to carry signals in one direction only (unidirectional), so one component 100 has a transmitter circuit 103 coupled to a signal source 110 labeled “normal path,” and one component 101 has a receiver circuit 104 coupled to a destination 111 labeled “normal path”. There are additional circuits present to permit periodic adjustment of the drive point and sample point in between periods of normal system operation. These adjustments compensate for changes in the system operating conditions.

[0042] The transmitter component includes a block 105 labeled “pattern”, which can consist of pattern storage, pattern generation circuitry or both, and which is used as a source of transmit calibration patterns for the first and second calibration sequences used according to the present invention. The first calibration sequence in other embodiments is provided by host software across the “normal” data path, while the second calibration sequence is provided by a “pattern” block in the transmitter, and vice versa. Generally, the pattern source can be the same or different, for the first exhaustive calibration sequence and second drift calibration sequence.

[0043] In the embodiment shown, a multiplexer block 106 labeled “mux,” implemented for example using a logical layer or physical layer switch, enables the transmit calibration pattern set to be driven onto the link by the transmitter circuit. The transmitter drive point can be adjusted by the block 107 labeled “adjust”. A sideband communication channel 113 is shown coupled between the component 101 and the component 100, by which the results of analysis of received calibration patterns at the component 101 are supplied to the adjust block 107 of the component 100.

[0044] The receiver component 101 includes a block 108 labeled “pattern”, which can consist of host software, pattern storage or pattern generation circuitry, and which is used as a source of expected patterns. A block 109 labeled “compare” enables the received pattern set to be compared to the expected pattern set, analyzes the result and causes an adjustment to be made to either the transmitter or receiver. The receiver sample point can be adjusted by the block 112 labeled “adjust”.

[0045] Fig. 4 shows two components 100, 101 connected with a unidirectional link 102, in which components of Fig. 3 are given like reference numerals. In the embodiment of Fig. 4, only the receiver sample point can be adjusted; the transmitter drive point remains fixed during system operation. Thus, there is no adjust block 107 in the component 100, nor is there a need for sideband communication channel 113 of Fig. 4.

[0046] Fig. 5 shows two components 100, 101 connected with a unidirectional link 102, in which components of Fig. 3 are given like reference numerals. In the embodiment of Fig. 5, only the transmitter drive point can be adjusted; the receiver sample point remains fixed during system operation. Thus, there is no adjust block 112 in the component 101 of Fig. 5.

[0047] In general, periodic timing calibration can be performed on all three examples, since timing variations due to condition drift can be compensated at either the transmitter end or the receiver end. In practice, it is cheaper to put the adjustment circuitry at only one end of the link, and not at both ends, so systems of Figs. 4 or 5 would have an advantage. Also, it should be noted that the system of Fig. 4 does not need to communicate information from the “compare” block 109 in the receiver component 101 back to the transmitter component 100, and thus might have implementation benefits over the system of Fig. 5. In practical applications, it is generally preferable to calibrate the receiver first and then use this setting to perform transmitter calibration. One method of breaking the chicken-and-egg problem is to load a known transmit pattern via a sideband interface to both sides of the channel. Having matching patterns on both sides, we can now perform an receive sample point calibration. Then, after the receiver offset is set to a reliable value, we can perform initial transmit drive point calibration.

[0048] Fig. 6 is a chart illustrating the measured values of a transmit drive timing point parameter using an exhaustive calibration sequence, as compared to the measured values of the same parameter using a short calibration sequence as discussed above. In Fig. 6, the conditions causing drift are plotted on the vertical axis, while the changes in the parameter are plotted on the horizontal axis. In the example illustrated, the condition causing drift is temperature, and the parameter being measured is the phase of the transmit drive point TX. As illustrated, the exhaustive calibration sequence computes edge values TXA and TXB on traces 200 and 201, so that the operation value TX as shown on trace 202 can be computed as a function of the edge values. In this case, the exhaustive calibration sequence utilizes a long calibration pattern including a pseudorandom bit sequence having a length  $2^N - 1$ , where N is at least 7, and N is 15 in one embodiment. Using a long pseudorandom bit sequence, the exhaustive center of the edge values TXA, TXB and the operation value TX are determined. In Fig. 6, the short calibration sequences are used to track drift, where the short sequences comprise a simple two-byte code or set of two byte (16 bit) codes, such as AAAA, 5555 or 0F0F (hexadecimal). As can be seen, the edge values TXAS and TXBS as measured using the short calibration sequences are plotted on traces 203 and 204. The operation value TXS, if it were based simply on the edge values TXAS and TXBS for the short calibration sequences would be on the trace 205. However, as shown in Fig. 6, the relationship between the condition on the vertical axis and the edge values TXAS, TXBS for the short calibration patterns along the horizontal axis has the same slope as the relationship between the condition and the operation value TX determined using the long calibration patterns. As shown in Fig. 6, if at calibration point C1, the operation value is established at TX, the edge value for TXAS determined using the short patterns has the value represented by point 206. At calibration point C2, the edge value TXAS drifts from point 206 on the horizontal axis by a drift value  $\Delta TXAS$ . Likewise, the operation value TX drifts by a drift value  $\Delta TX$ . To the extent the slope of trace 202 for the value TX is the same as, or otherwise correlates with, the slope of trace 203 for the value TXAS, the drift value  $\Delta TXAS$  can be applied directly to estimate the drift  $\Delta TX$ , and to update the operation value TX without executing an exhaustive calibration sequence. A similar relationship exists in the graph of Fig. 6, between TXAB and TX, and between TXS and

TX. That is the drift of TX is a function of the drift of TXAS alone, TXBS alone, or of a combination of TXAS and TXBS. Thus, by measuring changes in one or both of the edge values of the parameter using a short calibration pattern, and applying those changes to changing the operation value computed using a long calibration pattern, the operational center for the parameter can be adjusted to account for drift utilizing less resources of the communication channel.

[0049] Fig. 7 is a chart illustrating the measured values of a receive sample timing point parameter using an exhaustive calibration sequence, as compared to the measured values of the same parameter using a short calibration sequence. In Fig. 7, the conditions causing drift are plotted on the vertical axis, while the changes in the parameter are plotted on the horizontal axis. In the example illustrated, the condition causing drift is temperature, and the parameter being measured is the phase of the transmit drive point RX. As illustrated, the exhaustive calibration sequence computes edge values RXA and RXB on traces 210 and 211, so that the operation value RX as shown on trace 212 can be computed as a function of the edge values. In this case, the exhaustive calibration sequence utilizes a long calibration pattern including a pseudorandom bit sequence having a length  $2^N - 1$ , where N is at least 7, and N is 15 in one embodiment. Using a long pseudorandom bit sequence, the exhaustive center of the edge values RXA, RXB and the operation value RX are determined. In Fig. 7, the short calibration sequences are used to track drift, where the short sequences comprise a simple two-byte code or set of two byte (16 bit) codes, such as AAAA, 5555 or 0F0F (hexadecimal). As can be seen, the edge values RXAS and RXBS as measured using the short calibration sequences are plotted on traces 213 and 214. The operation value RXS, if it were based simply on the edge values RXAS and RXBS for the short calibration sequences would be on the trace 215.

However, as shown in Fig. 7, the relationship between the condition on the vertical axis and the edge values RXAS, RXBS for the short calibration patterns along the horizontal axis has the same slope as the relationship between the condition and the operation value RX determined using the long calibration patterns. As shown in Fig. 7, if at calibration point C1, the operation value is established at RX, the edge value for RXAS determined using the short patterns has the value represented by point 216. At calibration point C2, the edge value RXAS drifts from point 216 on the horizontal axis by a drift value

$\Delta RXAS$ . Likewise, the operation value  $RX$  drifts by an drift value  $\Delta RX$ . To the extent the slope of trace 212 for the value  $RX$  is the same as, or otherwise correlates with, the slope of trace 213 for the value  $RXAS$ , the drift value  $\Delta RXAS$  can be applied directly to estimate the drift  $\Delta RX$ , and to update the operation value  $RX$  without executing an exhaustive calibration sequence. A similar relationship exists in the graph of Fig. 7, between  $RXAB$  and  $RX$ , and between  $RXS$  and  $RX$ . That is the drift of  $RX$  is a function of the drift of  $RXAS$  alone,  $RXBS$  alone, or of a combination of  $RXAS$  and  $RXBS$ . Thus, by measuring changes in one or both of the edge values of the parameter using a short calibration pattern, and applying those changes to changing the operation value computed using a long calibration pattern, the operational center for the parameter can be adjusted to account for drift utilizing less resources of the communication channel.

[0050] The function of the drift in the edge value measured using the short calibration pattern, and the drift in the operation value using the long calibration pattern in the examples of Fig. 6 and Fig. 7, are linear, or approximately linear. In other examples, the function may be more complex, and/or the value measured using the short patterns may be different than those shown. However, so long as the drift in the value measured using the short calibration pattern, and the drift in the operation value using the long calibration pattern, are correlated, whether causally or not, then the values measured using the short calibration pattern can be applied for adjustment of the operation value of parameter.

[0051] Accordingly, the operation value of the parameter can be established, such as by using an exhaustive calibration routine based on long calibration patterns, during startup or initialization of the system. Adjustments in the operation value due to drift, however, can be made based upon shorter calibration sequences adapted to determine a drift value, such as those based on short calibration patterns. The shorter drift calibration sequences can be executed from time to time, such as on a set periodic basis, when signaled by an external monitor of the conditions of operation, or at times depending on other conditions of use of the communication channel, without utilizing as much of the resources of the communication channel as are needed for the more exhaustive calibration routine utilized less frequently, such as only during startup.

### **Calibration Steps for Transmitter for Unidirectional Link**

[0052] Fig. 8 shows calibration steps for the example system shown in Fig. 5.

(Step 300) Execute a first calibration sequence on initialization or other event to establish an operation value for the transmit drive point, based on long calibration patterns or more exhaustive routines.

5 (Step 301) Periodically, or upon an event indicating that a drift measurement is due, suspend transmit and receive operations that are in progress.

(Step 302) Change the drive point of the transmit component from the “TX” operation value (used for normal operations) to either the “TXAS” or “TXBS” edge value (used for drift calibration operations) in the “adjust” block. It may be necessary to impose a settling delay at this step to allow the new drive point to become stable.

10 (Step 303) Change “mux” block of the transmit component so that the “pattern” block input is enabled.

(Step 304) A pattern set is created in the “pattern” block of the transmit component and is transmitted onto the “link” using the TXAS or TXBS drive point.

15 (Step 305) The pattern set is received in the receive component. Note that the sample point of the receiver is fixed relative to the reference clock of the system and is not adjusted.

(Step 306) The received pattern set is compared in the “compare” block to the expected pattern set produced by the “pattern” block in the receive component. The two pattern sets will either match or not match. As a result of this comparison (and possibly other  
20 previous comparisons) a pass or fail determination will be made.

(Step 307) Adjust either the “TXAS” or “TXBS” edge value in the transmit component as a result of the pass or fail determination. The drift in the “TX” operation value in the transmit component is indicated by a function of the adjusted TXAS and TXBS values. The “TX” value is also adjusted based on that function. This adjustment may only be  
25 made after a calibration sequence including transmission of two or more of calibration patterns has been executed, in order to ensure some level of repeatability. Also, “low pass” filtering may be applied to measurements of drift or of the operation value, to prevent dithering of the operation value.

30 (Step 308) Change the drive point of the transmitter from the “TXAS” or “TXBS” edge value (used for calibration operations) to new “TX” operation value (used for normal operations) in the “adjust” block of the transmit component. It may be necessary to



impose a settling delay at this step to allow the new drive point to become stable.

(Step 309) Change “mux” block of the transmit component so that the “normal path” input is enabled.

(Step 310) Resume normal transmit and receive operations.

5 [0053] The first calibration sequence in one embodiment is an iteration just like that discussed above with respect to steps 301-310, with the exception that the calibration pattern for the first cycle is a long pattern, and the calibration pattern for the second calibration sequence is a short pattern. For example, a pseudorandom bit sequence having a length  $2^N - 1$ , with N equal to 7 or N equal to 15, can be used as the long  
10 calibration pattern for the first calibration sequence. In another example, the long calibration pattern is a set of short patterns intended to be an exhaustive set for conditions, such as inter-symbol interference patterns, of the communication channel. The short calibration pattern on the other hand, may be a simple two-byte code or set of two byte (16 bit) codes, such as AAAA, 5555 or 0F0F (hexadecimal).

15 [0054] In some embodiments, the exhaustive calibration sequence may use the same codes as the simpler calibration sequence, but apply shorter algorithms for computing adjustments. For example, the values may be adjusted without requiring repeatability, or requiring less repeatability, when measuring drift than when measuring the exhaustive operation value.

## 20 **Timing for Iteration Step for Transmit**

[0055] Fig. 9 includes the timing waveforms used by the exhaustive calibration steps to determine an operation value, similar to the calibration sequence shown in Fig. 8 for determining drift in a system like that of Fig. 5. These timing waveforms are similar to those from Fig. 1, except that the drive point is adjusted to straddle the sampling window  
25 of the receiver in order to track the edges of the valid window of the transmitter, and illustrate one example calibration procedure. During a drift calibration sequence, the edge values of Fig. 9, ( $t_{PHASET(TXA)}$  and  $t_{PHASER(TXB)}$ ) are replaced by correlated values determined using the short calibration sequences.

[0056] The “adjust” block in the transmit component maintains three values in  
30 storage: TXA, TX, and TXB. The TX value is the operation value used for normal operation. The TXA and TXB are the “edge” values, which track the left and right

extremes of the bit window of the transmitter. Typically, the TX value is initially derived from the average of the TXA and TXB values as determined using an exhaustive calibration sequence, but other relationships are possible.

[0057] As described above, drift in the TX value is determined by a function of the

5 TXAS and TXBS values determined during the calibration sequences that use short patterns. The function of TXAS and TXBS used depends on how the parameter drifts as measured by the exhaustive calibration patterns, as compared with how it drifts as measured by the short patterns. So long as this relationship correlates, then the short pattern technique for measuring drift is straight forward. In some systems, the drift in the  
10 short pattern measurements is very close to the drift in the long pattern measurements, so that any determined drift can be applied directly to adjust the operation value of the parameter.

[0058] The “adjust” block in the transmit component maintains three values in storage: TXA, TX, and TXB for the purposes of the exhaustive calibration sequence  
15 illustrated here. For drift calculation, TXAS and TXBS are also stored. The TX value is the operation value used for normal operation. The TXA and TXB are the “edge” values, which track the left and right extremes of the bit window of the transmitter. Typically, the TX value is initially derived from the average of the TXA and TXB values, but other relationships are possible. The TXA and TXB values can be maintained by the  
20 calibration operations, which from time to time, and periodically in some embodiments, interrupt normal operations for exhaustive calibration sequences. However, in embodiments of the present invention, storage of TXA and TXB may not be needed to track drift.

[0059] In Fig. 9, the position of the rising edge of CLK<sub>T</sub> has an offset of  $t_{PHASET}$   
25 relative to a fixed reference (typically a reference clock that is distributed to all components).

[0060] When the TX value is selected ( $t_{PHASET(TX)}$  in the middle trace 401 showing CLK<sub>T</sub> timing waveform) for operation, the rising edge 402 of CLK<sub>T</sub> causes the DATAT window 403 containing the value “a” to be aligned so that the DATAR signal (not shown  
30 but conceptually overlapping with the DATAT signal) at the receiving component is

aligned with the receiver clock, successfully received, and ideally centered on the receiver eye.

[0061] When the TXA value is selected ( $t_{PHASET(TXA)}$  in the top trace 405 showing CLK<sub>T</sub> timing waveform), the rising edge of CLK<sub>T</sub> is set to a time that causes the right edges of the DATAT window 406 (containing “a”) and the receiver set/hold window 410 (shaded) to coincide. The  $t_S$  setup time and  $t_H$  hold time surround the CLK<sub>R</sub> rising edge, together define the set/hold window 410 (not to be confused with the receiver eye of Fig. 2) in which the value of DATAR must be stable for reliable sampling around a given CLK<sub>R</sub> rising edge 404. Since the DATAT window, and the resulting DATAR window, are larger than this set/hold window 410, the transmitter has timing margin. However, in the case shown on trace 405 with the transmit clock rising edge at offset  $t_{PHASET(TXA)}$ , all the timing margin is on the left side of the transmitter eye for the set/hold window 410, adding delay after the  $t_Q$  timing parameter. There is essentially no margin for the  $t_V$  timing parameter in the trace 405, so that the offset defines the left edge of the calibration window.

[0062] The calibration process for TXA will compare the received pattern set to the expected pattern set, and determine if they match. If they match (pass) then the TXA value will be decremented (the  $t_{PHASET(TXA)}$  offset becomes smaller shifting the transmit window 406 to the left in Fig. 9) or otherwise adjusted, so there is less margin for the  $t_V$  timing parameter relative to the receiver window 410. If they do not match (fail) then the TXA value will be incremented (the  $t_{PHASET(TXA)}$  offset becomes larger shifting the transmit window 406 to the right in Fig. 9) or otherwise adjusted, so there is more margin for the  $t_V$  timing parameter.

[0063] As mentioned earlier, the results of a sequence including transmission of two or more calibration patterns may be accumulated before the TXA value is adjusted. This would improve the repeatability of the calibration process. For example, the calibration pattern could be repeated “N” times with the number of passes accumulated in a storage element. If all N passes match, then the TXA value is decremented. If any of the N passes do not match, then the TXA value is determined to have reached the edge of the window and is incremented. In another alternative, after the Nth pattern, the TXA value

could be incremented if there are fewer than  $N/2$  (or some other threshold number) passes, and decremented if there are  $N/2$  or more passes.

[0064] When TXA is updated, the TX value will also be updated. In this example, the TX value will be updated by half the amount used to update TXA, since TX is the average of the TXA and TXB values. If TX has a different relationship to TXA and TXB, the TX update value will be different. Note that in some embodiments, the TX value will need slightly greater precision than the TXA and TXB values to prevent round-off error. In alternate embodiments, the TX value can be updated after pass/fail results of TXA and TXB values have been determined. In some cases, these results may cancel and produce no change to the optimal TX value. In other cases these results may be accumulated and the accumulated results used to determine an appropriate adjustment of the TX setting. According to this embodiment, greater precision of the TX setting relative to the TXA and TXB settings may not be required.

[0065] When the TXB value is selected ( $t_{\text{PHASER(TXB)}}$  in the bottom trace 407 showing a CLK\_T timing waveform) for calibration, the rising edge of CLK\_T is set to a time that causes the left edge of the transmitter valid window 408 (containing "a") and the receiver set/hold window 410 (shaded) to coincide. In this case with the transmit clock rising edge at  $t_{\text{PHASER(TXB)}}$ , all the timing margin is on the right side of the transmit window 408, providing more room than required by the  $t_v$  timing parameter. This means that there will be essentially no margin for the  $t_q$  timing parameter on the left side of the window 408, defining the right edge of the calibration window.

[0066] The calibration process will compare the received pattern set to the expected pattern set, and determine if they match. If they match (pass) then the TXB value will be incremented (the offset becomes larger) or otherwise adjusted, so there is less margin for the  $t_q$  timing parameter. If they do not match (fail) then the TXB value will be decremented (the offset becomes smaller) or otherwise adjusted, so there is more margin for the  $t_q$  timing parameter.

[0067] As mentioned earlier, the results of transmission of two or more calibration patterns may be accumulated before the TXB value is adjusted. For example, transmission of the patterns could be repeated "N" times with the number of passes accumulated in a storage element. After the Nth sequence the TXB value could be

decremented if there are fewer than  $N/2$  passes and incremented if there are  $N/2$  or more passes. This would improve the repeatability of the calibration process.

[0068] When TXB is updated, the TX value will also be updated. In this example, the TX value will be updated by half the amount used to update TXB, since TX is the average of the TXA and TXB values. If TX has a different relationship to TXA and TXB, the TX update value will be different. Note that the TX value will need slightly greater precision than the TXA and TXB values if it is desired to prevent round-off error.

[0069] Determination of drift may be made using a similar process, although with different calibration patterns, as explained above.

## Calibration Steps for Receiver for Unidirectional Link

[0070] Fig. 10 shows, for a system like the example from Fig. 4, the steps needed to perform a timing calibration update for a receiver. Note that only steps (Block 502), (Block 507), and (Block 508) are different relative to the steps in Fig. 8.

(Step 500) Execute a first calibration sequence on initialization or other event to establish an operation value "RX" for the receive sample point, based on long calibration patterns that are intended to be exhaustive.

(Step 501) Periodically, or upon an event indicating that a drift measurement is due, suspend normal transmit and receive operations that are in progress.

(Step 502) Change the sample point of the receive component from the "RX" operation value (used for normal operations) to either the "RXAS" or "RXBS" edge value (used for calibration operations) in the "adjust" block. It may be necessary to impose a settling delay at this step to allow the new sample point to become stable.

(Step 503) Change "mux" block of the transmit component so that the "pattern" block input is enabled.

(Step 504) A pattern set is created in the "pattern" block of the transmit component and is transmitted onto the "link" using the TX drive point. The pattern set includes one or more short codes, for example one or more 2-byte codes.

(Step 505) The pattern set is received in the receive component.

(Step 506) The received pattern set is compared in the "compare" block to the expected pattern set produced by the "pattern" block in the receive component. The two pattern sets will either match or not match. As a result of this comparison (and possibly other

previous comparisons) a pass or fail determination will be made.

(Step 507) Adjust either the “RXAS” or “RXBS” edge value in the receive component as a result of the pass or fail determination. The drift in the “RX” operation value in the transmit component is indicated by a function of the adjusted RXAS and RXBS values.

5 The “RX” value is also adjusted based on that function. This adjustment may only be made after two or more of these calibration sequences have been executed, in order to ensure some level of repeatability.

(Step 508) Change the sample point of the receiver from the “RXAS” or “RXBS” edge value (used for calibration operations) to “RX” operation value (used for normal  
10 operations) in the “adjust” block of the receive component. It may be necessary to impose a settling delay at this step to allow the new drive point to become stable.

(Step 509) Change “mux” block of the transmit component so that the “normal path” input is enabled.

(Step 510) Resume normal transmit and receive operations.

15 **[0071]** The first calibration sequence in one embodiment is an iteration just like that discussed above with respect to steps 501-510, with the exception that the calibration pattern for the first cycle is a long pattern, and the calibration pattern for the second calibration sequence is a short pattern. For example, a pseudorandom bit sequence having a length  $2^N - 1$ , with N equal to 7 or N equal to 15, can be used as the long  
20 calibration pattern for the first calibration sequence. In another example, the long calibration pattern is a set of short patterns intended to be an exhaustive set for conditions, such as inter-symbol interference patterns, of the communication channel. The short calibration pattern on the other hand, may be a simple two-byte code or set of two-byte codes, such as AAAA, 5555 or 0F0F.

## 25 **Timing for Iteration Step for Receive**

**[0072]** Fig. 11 shows the exhaustive calibration steps to determine an operation value, similar to the calibration sequence shown in Fig. 8 for determining drift of the operation value, in a system like that of Fig. 4. These timing waveforms are similar to those from Fig. 1, except that the sampling point is adjusted within the bit window in  
30 order to track the edges of the window. During a drift calibration sequence, the edge

values of Fig. 11, ( $t_{PHASET(RXA)}$  and  $t_{PHASER(RXB)}$ ) are replaced by correlated values determined using the short calibration sequences.

[0073] The “adjust” block in the receive component maintains three values in storage: RXA, RX, and RXB for the purpose of the exhaustive calibration sequence. The RX value is the operation value used for normal operation. The RXA and RXB are the “edge” values, which track the left and right extremes of the bit window. The RX value is derived from the first exhaustive calibration sequence, and based on the average of the RXA and RXB values determined in the first calibration sequence, but other relationships are possible. The RXAS and RXBS values are determined in the second calibration sequence using the short calibration patterns, and maintained by the calibration operations which periodically interrupt normal operations. Changes in the RXAS and RXBS values, or in only one of them, are used to determine a drift in the RX values. The RX value is updated based on that drift.

[0074] In the timing diagrams, the position of the rising edge of CLKR has an offset of  $t_{PHASER}$  relative to a fixed reference (not shown, typically a reference clock that is distributed to all components). This offset is determined by the RXA, RX, and RXB values that are stored.

[0075] When the RX value is selected ( $t_{PHASER(RX)}$  in the middle trace 601 showing a CLKR timing waveform) for use in receiving data, the rising edge 602 of CLKR is approximately centered in the receiver eye of the DATAR signal containing the value “a”. The DATAR signal is the DATAT signal transmitted at the transmitter after propagation across the link, and can be conceptually considered to be the same width as DATAT as shown in Fig. 11. The receiver eye is shown in Fig. 2. The  $t_s$  setup time is the minimum time before the clock CLKR rising edge which must be within the DATAR window 603, and the  $t_H$  hold time is the minimum time after the clock CLKR rising edge that must be within the DATAR window 603, together defining the set/hold window 604 (not to be confused with the receiver eye of Fig. 2) in which the value of DATAR must be stable for reliable sampling around a given CLKR rising edge. Since the valid window 604 of the DATAR signal is larger than this set/hold window 604, the receiver has timing margin in both directions.

[0076] When the RXA value is selected ( $t_{\text{PHASER(RXA)}}$  in the top trace 605 showing a CLKR timing waveform), the rising edge of CLKR is approximately a time  $t_s$  later than the left edge (the earliest time) of the DATAR window 603 containing the value “a”. In this case, the CLKR rising edge is on the left edge of the receiver eye, and all the timing margin is on the right side of the set/hold window 604, providing more room than is required by the  $t_H$  timing parameter. This means that there will be essentially no margin for the  $t_s$  timing parameter, defining the left edge of the calibration window.

[0077] The calibration process will compare the received pattern set to the expected pattern set, and determine if they match. If they match (pass) then the RXA value will be decremented (the offset becomes smaller) or otherwise adjusted, so there is less margin for the  $t_s$  timing parameter. If they do not match (fail) then the RXA value will be incremented (the offset becomes larger) or otherwise adjusted, so there is more margin for the  $t_s$  timing parameter.

[0078] As mentioned earlier, the results of transmission and reception of two or more calibration patterns may be accumulated before the RXA value is adjusted. For example, the patterns could be repeated “N” times with the number of passes accumulated in a storage element. After the Nth sequence the RXA value could be incremented if there are fewer than  $N/2$  passes and decremented if there are  $N/2$  or more passes. This would improve the repeatability of the calibration process.

[0079] When RXA is updated, the RX value will also be updated. In this example, the RX value will be updated by half the amount used to update RXA, since RX is the average of the RXA and RXB values. If RX has a different relationship to RXA and RXB, the RX update value will be different. Note that in some embodiments, the RX value will need slightly greater precision than the RXA and RXB values to prevent round-off error. In alternate embodiments, the RX value can be updated after pass/fail results of RXA and RXB values have been determined. In some cases, these results may cancel and produce no change to the optimal RX value. In other cases these results may be accumulated and the accumulated results used to determine an appropriate adjustment of the RX setting. According to this embodiment, greater precision of the RX setting relative to the RXA and RXB settings may not be required.



[0080] When the RXB value is selected ( $t_{\text{PHASER}(\text{RXB})}$  in the bottom trace 606 showing a CLKR timing waveform), the rising edge of CLKR is approximately a time  $t_H$  earlier than the right edge (the latest time) of the DATAR window 603 containing the value “a”. In this case, the CLKR rising edge is on the right edge of the receiver eye, and all the  
 5 timing margin is on the left side of the window 604, providing more room that required by the  $t_S$  timing parameter. This means that there will be essentially no margin for the  $t_H$  timing parameter, defining the right edge of the calibration window.

[0081] The calibration process will compare the received pattern set to the expected pattern set, and determine if they match. If they match (pass) then the RXB value will be  
 10 incremented (the offset becomes larger) or otherwise adjusted, so there is less margin for the  $t_H$  timing parameter. If they do not match (fail) then the RXB value will be decremented (the offset becomes smaller) or otherwise adjusted, so there is more margin for the  $t_H$  timing parameter.

[0082] As mentioned earlier, the results of transmission and reception of two or more  
 15 calibration patterns may be accumulated before the RXB value is adjusted. For example, the sequence could be repeated “N” times with the number of passes accumulated in a storage element. After the Nth sequence the RXB value could be decremented if there are fewer than  $N/2$  passes and incremented if there are  $N/2$  or more passes. This would improve the repeatability of the calibration process.

[0083] When RXB is updated, the RX value will also be updated. In this example,  
 20 the RX value will be updated by half the amount used to update RXB, since RX is the average of the RXA and RXB values. If RX has a different relationship to RXA and RXB, the RX update value will be different. Note that the RX value will need slightly greater precision than the RXA and RXB values if it is desired to prevent round-off error.

[0084] Determination of drift may be made using a similar process, although with  
 25 different calibration patterns, as explained above.

### **Bidirectional Link Alternatives**

[0085] Fig. 12 shows an example of a bidirectional link. In this case, component A  
 30 (700) and component B (701) each contain a transmitter and receiver connected to the link, so that information may be sent either from A to B or from B to A. The elements of

the unidirectional example in Fig. 3 are replicated (two copies) to give the bidirectional example in Fig. 11. Fig. 11 shows two bidirectional components 700, 701 connected with an interconnection medium referred to as Link 702. Normal path 710 acts as a source of data signals for normal operation of component 700 during transmit operations. Normal path 731 acts as a destination of data signals for component 700 during normal receive operations. Likewise, normal path 730 acts as a source of data signals for normal operation of component 701 during transmit operations. Normal path 711 acts as a destination of data signals for component 701 during normal receive operations.

[0086] The first bidirectional component includes a block 705 labeled “pattern”, which can consist of pattern storage, pattern generation circuitry, or both, and which is used as a source of transmit calibration patterns for both the exhaustive calibration sequences and the calibration sequences used to measure drift. A multiplexer block 706 labeled “mux,” implemented for example using a logical layer or physical layer switch, enables the transmit calibration pattern set to be driven onto the link by the transmitter circuit 703. The transmitter drive point can be adjusted by the block 707 labeled “adjust”. A sideband communication channel 713 is shown coupled between the component 701 and the component 700, by which the results of analysis of received calibration patterns at the component 701 are supplied to the adjust block 707 of the component 700. Component 700 also has support for calibrating receiver 724, including a block 728 labeled “pattern”, which can consist of pattern storage, pattern generation circuitry, or both, and which is used as a source of expected patterns for comparison with received patterns. A block 729 labeled “compare” enables the received pattern set to be compared to the expected pattern set, and causes an adjustment to be made to either the transmitter or receiver. The receiver sample point can be adjusted by the block 732 labeled “adjust”.

[0087] The second bidirectional component 701 includes complementary elements supporting transmitter 723 and receiver 704. For the receiver operations, a block 708 labeled “pattern”, which can consist of pattern storage, pattern generation circuitry, or both, is used as a source of expected patterns. A block 709 labeled “compare” enables the received pattern set to be compared to the expected pattern set, and causes an adjustment to be made to either the transmitter or receiver. The receiver sample point can

be adjusted by the block 712 labeled “adjust”. The second bidirectional component 701 supports transmission operations, with elements including a block 725 labeled “pattern”, which can consist of pattern storage or pattern generation circuitry, and which is used as a source of transmit calibration patterns. A multiplexer block 726 labeled “mux,”

5 implemented for example using a logical layer or physical layer switch, enables the transmit calibration pattern set to be driven onto the link by the transmitter circuit 723.

The transmitter drive point can be adjusted by the block 727 labeled “adjust”. A sideband communication channel 733 is shown coupled between the component 700 and the component 701, by which the results of analysis of received calibration patterns at the  
10 component 700 are supplied to the adjust block 727 of the component 701.

**[0088]** Fig. 12 allows both receive sample points and both transmit drive points to be adjusted. However, the benefit of adjustable timing will be realized if there is only one adjustable element in each direction.

**[0089]** The example of Fig. 13 (using the same reference numerals as Fig. 12) shows  
15 an example in which only the receiver sample points are adjustable. Thus, elements 707 and 727 of Fig. 12 are not included in this embodiment. This is equivalent to two copies of the elements of example in Fig. 4.

**[0090]** The example of Fig. 14 (using the same reference numerals as Fig. 12) shows an example in which only the transmitter drive points are adjustable. Thus, elements 712  
20 and 732 of Fig. 12 are not included in this embodiment. This is equivalent to two copies of the elements of example in Fig. 5.

**[0091]** Example of Fig. 15 (using the same reference numerals as Fig. 12) shows an example in which the receiver sample point and transmitter drive point of the first bidirectional component 700 are adjustable. Thus, elements 712, 708, 709, 727, 726, 725  
25 are not included in this embodiment. A storage block 750 is added between the receiver and a “mux” block 751. The “mux” block 751 is used to select between a normal source of signals 730 and the storage block 750. Also, the compare block 752 is used for analysis of both transmit and receive calibration operations, and is coupled to both the adjust block 707 for the transmitter, and adjust block 732 for the receiver. This  
30 alternative is important because all the adjustment information can be kept within one component, eliminating the need for sideband signals for the calibration process. If

component 701 were particularly cost sensitive, this could also be a benefit, since only one of the components must bear the cost of the adjustment circuitry.

### **Calibration Steps for Transmitter for Bidirectional Link**

[0092] The calibration steps for bidirectional examples in Figs. 12, 13 and 14 will be essentially identical to the calibration steps already discussed for unidirectional examples in Figs. 4 and 5. However, the asymmetry in bidirectional example of Fig. 15 will introduce some additional calibration steps, and will receive further discussion.

[0093] Fig. 16 includes the steps needed to perform a timing calibration update for a system like that of Fig. 15.

(Step 800) Execute a first calibration sequence on initialization or other event to establish an operation value "TX" for the transmit drive point, based on long calibration patterns that are intended to be exhaustive, or to provide results usable in the intended conditions for the system.

(Step 801) Periodically, or upon an event indicating that a drift measurement is due, suspend normal transmit and receive operations that are in progress.

(Step 802) Change the drive point of the transmit component (A) from the "TX" operation value (used for normal operations) to either the "TXAS" or "TXBS" edge value (used for calibration operations) in the "adjust" block. It may be necessary to impose a settling delay at this step to allow the new drive point to become stable.

(Step 803) Change "mux" block of the transmit component (A) so that the "pattern" block input is enabled.

(Step 804) A pattern set is created in the "pattern" block of the transmit component (A) and is transmitted onto the "link" using the TXAS or TXBS drive point.

(Step 805) The pattern set is received in the receive component (B). Note that the sample point of the receiver is fixed relative to the reference clock of the system and is not adjustable. The received pattern set is held in the "storage" block in component B.

(Step 806) The "mux" block input connected to the "storage" block in component B is enabled. The pattern set is re-transmitted onto the link by component B.

(Step 807) The pattern set is received by component B from the link.

(Step 808) The received pattern set is compared in the "compare" block to the expected pattern set produced by the "pattern" block in the receive component (A). The two

pattern sets will either match or not match. As a result of this comparison (and possibly other previous comparisons) a pass or fail determination will be made.

(Step 809) Adjust either the “TXAS” or “TXBS” edge value in the transmit component (A) as a result of the pass or fail determination to determine a drift value. The “TX”

5 operation value in the transmit component (A) is also adjusted as a result of the determined drift value. This adjustment may only be made after two or more of these calibration sequences have been executed, in order to ensure some level of repeatability.

(Step 810) Change the drive point of the transmitter from the “TXAS” or “TXBS” edge value (used for calibration operations) to “TX” operation value (used for normal

10 operations) in the “adjust” block of the transmit component (A). It may be necessary to impose a settling delay at this step to allow the new drive point to become stable.

(Step 811) Change “mux” block of the transmit component (A) so that the “normal path” input is enabled.

(Step 812) Resume normal transmit and receive operations.

## 15 **Calibration Steps for Receiver for Bidirectional Link**

**[0094]** The calibration steps for the bidirectional examples of Figs. 12, 13, and 14 will be essentially identical to the calibration steps already discussed for the unidirectional examples of Figs. 4 and 5. However, the asymmetry in bidirectional example of Fig. 15 will introduce some additional calibration steps, and will receive  
20 further discussion.

**[0095]** Fig. 17 includes the steps needed to perform a receiver timing calibration update for a system like that of Fig. 15.

(Step 900) Execute a first calibration sequence on initialization or other event to establish an operation value “RX” for the receive sample point, based on long calibration patterns.

25 (Step 901) Periodically, or upon an event indicating that a drift measurement is due, suspend normal transmit and receive operations that are in progress.

(Step 902) Change the sample point of the receive component (A) from the “RX” operation value (used for normal operations) to either the “RXAS” or “RXBS” edge value (used for calibration operations) in the “adjust” block. It may be necessary to

30 impose a settling delay at this step to allow the new drive point to become stable.

(Step 903) Change “mux” block of the transmit component (A) so that the “pattern”

block input is enabled.

(Step 904) A pattern set is created in the “pattern” block of the transmit component (A) and is transmitted onto the “link”. The normal transmit drive point is used.

(Step 905) The pattern set is received in the receive component (B). Note that the sample point of the receiver in component (B) is fixed relative to the reference clock of the system and is not adjustable. The received pattern set is held in the “storage” block in component B.

(Step 906) The “mux” block input connected to the “storage” block in component B is enabled. The pattern set is re-transmitted onto the link by component B.

(Step 907) The pattern set is received by component A from the link using either the RXAS or RXBS value to determine the receiver sample point.

(Step 908) The received pattern set is compared in the “compare” block to the expected pattern set produced by the “pattern” block in the receive component (A). The two pattern sets will either match or not match. As a result of this comparison (and possibly other previous comparisons) a pass or fail determination will be made.

(Step 909) Adjust either the “RXAS” or “RXBS” edge value in the receive component (A) as a result of the pass or fail determination to determine a drift value. The “RX” operation value in the receive component (A) is also adjusted based on this drift value. This adjustment may only be made after two or more of these calibration sequences have been executed, in order to ensure some level of repeatability.

(Step 910) Change the sample point of the receiver from the “RXAS” or “RXBS” edge value (used for calibration operations) to “RX” operation value (used for normal operations) in the “adjust” block of the receive component (A). It may be necessary to impose a settling delay at this step to allow the new drive point to become stable.

(Step 911) Change “mux” block of the transmit component (A) so that the “normal path” input is enabled.

(Step 912) Resume normal transmit and receive operations.

**[0096]** Fig. 18 illustrates an example like that of Fig. 15, with the exception that the point to point bidirectional link of Fig. 15 is replaced with a multidrop link, coupling component 1000 to a plurality of components 1051, 1052. The multidrop link configuration can be applied in other configurations. In the representative example

shown in Fig. 18, a first bidirectional component 1000 and a plurality of other bidirectional components 1051, 1052 are connected in a point to multi-point configuration, or multipoint to multipoint configuration, with an interconnection medium referred to as Link 1002. Normal path 1010 acts as a source of data signals for normal operation of component 1000 during transmit operations. Normal path 1031 acts as a destination of data signals for component 1000, during normal receive operations. The calibration operations are interleaved, and re-ordered, in this embodiment with normal communications, as described above to improve throughput and utilization of the communication medium.

[0097] The first bidirectional component 1000 includes a block 1005 labeled “pattern”, which can consist of pattern storage or pattern generation circuitry, and which is used as a source of transmit calibration patterns. A multiplexer block 1006 labeled “mux,” implemented for example using a logical layer or physical layer switch, enables the transmit calibration pattern set to be driven onto the link by the transmitter circuit 1003. The transmitter drive point can be adjusted by the block 1007 labeled “adjust”. In this embodiment, the adjust block 1007 includes storage for multiple parameter sets which are applied depending on the one of the other components 1051, 1052, ... on the link to which the transmission is being sent. Component 1000 also has support for calibrating receiver 1024, including a block 1028 labeled “pattern”, which can consist of pattern storage or pattern generation circuitry, and which is used as a source of expected patterns for comparison with received patterns. A block 1029 labeled “compare” enables the received pattern set to be compared to the expected pattern set, and causes an adjustment to be made to either the transmitter or receiver. The receiver sample point can be adjusted by the block 1032 labeled “adjust”. In this embodiment, the adjust block 1007 includes storage for multiple parameter sets which are applied depending on the one of the other components 1051, 1052, ... on the link from which the communication is being received. In the first component 1000, the compare block 1029 is used for analysis of both transmit and receive calibration operations, and is coupled to both the adjust block 1007 for the transmitter, and adjust block 1032 for the receiver. In the example of Fig. 18, the receiver sample point and transmitter drive point of the first bidirectional component 1000 are adjustable. The other components 1051, 1052, ... are implemented

as described with reference to Fig. 15 without adjustment resources, in this example, and not described here. In alternative embodiments, the components 1051, 1052, ... on the link may be provided with adjustment and calibration resources, as described for other embodiments above.

5    **[0098]**    The embodiments described above involve calibration of timing parameters. Other embodiments of the invention are applied to calibration of other parameters of the communication channel, including voltage levels for drivers and comparators, resistance values such as link termination resistances, driver strength, adaptive equalization coefficients, noise cancellation coefficients, parameters that cause overshoot and  
10   undershoot of signals such as driver switching power or speed, and so on. These parameters are reflected in drifting flight times, output delays for transmitters, receiver mismatches and input delays. Also, drift occurs due to spread spectrum clocking.

**[0099]**    The calibration sequences used for tracking drift are executed from time to time. In some embodiments of the invention, the calibration sequences for tracking drift  
15   are initiated in response to a timer, or upon sensing a change in a condition that could cause drift such as ambient temperature, temperature of a component or printed circuit board, power supply voltage variations, system state changes like switches between low and higher speed or power modes of operation, host system or system management commands, and so on. The timing for the calibration sequences is selected in preferred  
20   systems depending on time constants expected for the drifting conditions likely to have an effect on the parameters being calibrated.

**[00100]**   According to the present invention, communication systems are provided that support periodic calibration to track drift, while conserving resources of the communication channels being calibrated. The invention reduces the hardware and  
25   storage requirements needed for calibration, reduces the impact on bandwidth and throughput across the communication channel, and maintains accuracy of the operation values of the parameters being calibrated while using less of the resources of the communication channel. The calibration sequences tracking drift can be run more often, because overhead is smaller, than prior art schemes. Also, the techniques of the present  
30   invention improve overall latency characteristics of the communication channel in changing operating conditions.



**[00101]** While the present invention is disclosed by reference to the preferred embodiments and examples detailed above, it is to be understood that these examples are intended in an illustrative rather than in a limiting sense. It is contemplated that modifications and combinations will readily occur to those skilled in the art, which  
5 modifications and combinations will be within the spirit of the invention and the scope of the following claims. What is claimed is: